ACS Medicinal Chemistry Letters

Viewpoint

Cite This: ACS Med. Chem. Lett. 2018, 9, 1150−1152

pubs.acs.org/acsmedchemlett

Artificial Intelligence in Drug Design—The Storm Before the Calm?

Allan M. Jordan*

Drug Discovery Unit, Cancer Research UK Manchester Institute, The University of Manchester, Alderley Park, Macclesfield, SK10 4TG, U.K.

ABSTRACT: Artificial intelligence in drug design in experiencing a wave of excitement not seen since the emergence of computational chemistry in the late 1980s and early 1990s. Apparently failing to learn the lessons of recent history, we are promised imminent and pervasive solutions to the ills of drug design and significant increases in productivity as we seek to deliver innovative new therapeutics. However, do significant issues remain to be answered before AI enters the day-to-day toolbox of the practicing medicinal chemist?

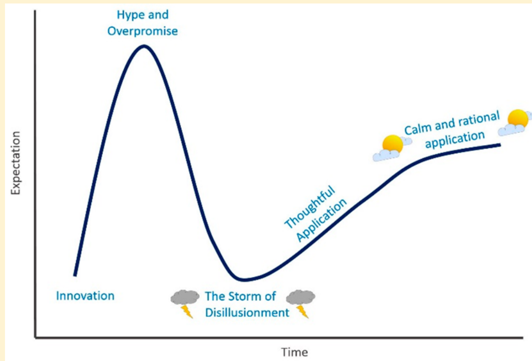

Back in the late 1980s and early 1990s, the advent of (for the time) high powered computational facilities heralded a new era in computational drug design. Here, for the first time, was a technology that could be used to interpret, de novo, emerging three-dimensional data on protein drug targets and deliver novel new molecules for synthesis. Computer-aided drug design was feted with the ability to design unique and bespoke molecular scaffolds with unprecedented levels of potency and selectivity. Small molecules would be designed to interact with key binding site residues and, in doing so, render largely obsolete the irrational and biassed designs of the medicinal chemist. No longer would there be a requirement to design hundreds, or thousands, of molecules. Here was a technology that might deliver "the one"—that single molecule designed to bind specifically to any given drug target with exquisite potency and selectivity, fast-tracking delivery of molecules though to clinical evaluation.

Fast forward three decades, and the reality is somewhat different. As is commonplace with new technologies, computational chemistry has experienced all facets of the Gartner Hype Curve, passing through overhyped expectation, failure to deliver on promises made, despair, and finally through to a useful tool, employed alongside other toolbox components available to the medicinal chemist. Without doubt, structure-guided drug design has accelerated the path to clinic for many therapeutic agents, including inhibitors of kinases and chaperone proteins, GPCR antagonists, and compounds destabilizing protein−protein interactions. Analysis of drug target structure has expedited unprecedented levels of selectivity to be engineered against seemingly identical binding sites. Understanding of metabolic enzymes, at the atomic level, has allowed the prediction and optimization of the molecular stabilities of emerging drug candidates. Predictive models of pharmacokinetic and physicochemical properties have facilitated the accelerated triage of vast libraries of hypothetical compounds, and QSAR modeling has enabled the application of accrued data to predict and prioritize the biochemical and pharmacokinetic properties of possible new compounds for synthesis.

Yet, despite ever-increasing computational power, better force field models, and an exponential growth of real data upon which to base models, the original dream of computational chemistry has yet to be realized. Thus, far, the practicing medicinal chemist cannot simply import their protein structure of choice into an educated system, drawing upon three decades of knowledge and learning, and have the system predict, with certainty, the one key molecule with the sublime balance of properties required to be a drug suitable for human clinical trials.

We accept these limitations, work alongside them, and have learned to apply the plethora of computational chemistry tools to our projects in appropriate ways to accelerate our studies while accepting and working within their shortfalls to deliver insight, direction, and focus. We use computational models in a realistic manner as an adjunct, rather than a panacea of prediction.

Yet, despite the checkered past of computational chemistry, we seem unwilling to learn from our experiences and heartache. Once again, the advent of new technologies to solve the problems of the medicinal chemist promises much but stands poised to underdeliver on the hype and hyperbole attributed to it. Indeed, the current rhetoric around artificial intelligence (AI) strongly echoes that which surrounded structure-guided drug design barely a few decades ago, and it is not unreasonable to expect similar levels of disappointment and distrust to arise if the technology fails to deliver upon the significant deliverables that have been promised. At a variety of recent meetings, AI has been heralded as a game-changer in

Published: November 2, 2018

© 2018 American Chemical Society

DOI: 10.1021/acsmedchemlett.8b00500
ACS Med. Chem. Lett. 2018, 9, 1150−1152

drug discovery; and it will, undoubtedly, have significant impact. However, it is the author's opinion that there remain significant challenges and hurdles to be overcome before this technology can truly deliver benefit to patients. At risk of being a lone voice in the wilderness and being branded an AI Luddite,[1] significant challenges surely loom on the horizon before this technology can be truly game-changing in our day-to-day activities.

First, and most critically, AI simply is not intelligent in the truest sense, at least not yet. Rather, the algorithms upon which the present systems are based rely heavily upon machine learning and pattern recognition. True intelligence in drug design suggests the ability to develop hypotheses and irrational suggestions based on creativity, insight, and gut feeling. In the present incarnations, these systems draw together vast data sets with mind-boggling efficiency to deliver reasoned suggestions based on historical precedent, but are not yet applying true intelligent creativity to drive projects forward. To use a retail analogy, this suggestive algorithm is much more akin to the purchasing model of major online retailers ("Chemists who made this, also made this...") based on big data learning, as opposed to a truly creative and imaginative suggestion.

While this approach has clear value for those targets where such precedent exists (e.g., at the protein family or target class level), few data sets seem to exist with demonstrable utility in less well studied cases. As we continue to move toward novel therapeutic paradigms, it will be revealing to see how well these data aggregation approaches can add value to the hitherto undocumented problems associated with a growing list of new modalities, such as PROTACs, interference with microRNAs, chaperones of misfolded proteins, alteration of mRNA splicing, activators of read-through transcription, nanobody design and delivery, and other future points of therapeutic intervention.

One potential drawback of this machine learning and pattern matching methodology is that it has the propensity to become self-reinforcing in certain areas. Take, for example, the emerging platforms for reaction planning. Learning from vast tranches of literature, and understanding the specific limitations of reactions in terms of functional group tolerance, protecting group requirements, and preferable reagents, and combining these findings with a scope that is orders of magnitude more comprehensive than even the most well-read synthetic chemist, these systems can suggest truly innovative routes to new molecules. Importantly, such systems are demonstrating validity in terms of reducing these ideas to practice.[2] However, in suggesting best routes, such systems also often prioritize those routes based on frequency of utilization, i.e., those most used in analogue generation, such as Pd-mediated couplings and amide formations. In doing so, it is not unreasonable to suggest that over time, the utilization of these reactions increases ever further, making them more likely to be suggested by, and therefore potentially reducing the desired creativity and power of the AI systems to help us avoid over-reliance on certain reaction types.[3−5]

Let us not forget that access to such systems requires either significant internal investment or the ability to "buy in" technology at significant cost. One of the critical benefits of the present incarnation of computational chemistry, after decades of development, is the accessibility of the systems to considerable numbers of chemists, from industry to academia and not-for-profit drug discovery laboratories. This, combined with intuitive interfaces, helps advance our projects through multiple parallel interactions with the technology and the triage of the resultant ideas. Whilst some efforts are underway to make AI tools more widely accessible to the general community,[6] they often do not have the validation of more developed platforms and their utility remains to be fully evaluated.

Derek Lowe (of "In the Pipeline" fame[7]), discussing AI, has expressed the sentiment that "AI will not be the end of medicinal chemist, but it will be the end of medicinal chemists who do not use AI".[8]

While I believe this will become true in the longer term, in the near term, I believe the AI community has a significant part to play to make this happen. To be truly useful, AI in medicinal chemistry needs to learn from the lessons of the rise, fall, and rise again of computational chemistry, avoiding the catastrophic fall from trust to mistrust, before rising phoenix-like from the ashes as a credible and useful tool. AI must evolve as a platform to become truly applicable to novel problems, with limited prior data sets. It needs to apply lateral decision making to make nonobvious and illogical connections across incomplete data sets, in the way a trained medicinal chemist does routinely in their day-to-day role. It needs to apply tacit knowledge to unforeseen problems in a meaningful way. It is imperative that the AI systems suggest credible, synthetically accessible molecules with rational and understandable arguments for reduction to practice, and in doing so, be faster, more efficient , and more creative than the chemist themselves; finally, do all of this in a way that is cost-effective and accessible, such that it becomes a standard tool for idea exploration on the desktop of every chemist—suggesting, prioritizing, and focusing thinking in a visual and meaningful manner in just a few mouse clicks.

Artificial Intelligence in drug design is currently riding a storm of interest. Similar to many medicinal chemists dedicated to delivering benefit to patients, I look forward to the coming calm, when we can truly use this technology to benefit those whom we seek to help. For those who the present storm of enthusiasm and expectation offers hope, be aware that it may deliver significant disappointment first. For those whom this new technological investment may offer a significant step forward in the delivery of new and much needed therapeutics, I believe we owe it to them to ensure we learn from our prior experiences and make sure these technologies are applied appropriately, and realistically, to the drug discovery challenges we face in our day-to-day roles.

AUTHOR INFORMATION

Corresponding Author

*E-mail: allan.jordan@cruk.manchester.ac.uk.

ORCID

Allan M. Jordan: 0000-0003-3449-3993

Notes

Views expressed in this editorial are those of the author and not necessarily the views of the ACS.

The opinions expressed in this viewpoint are entirely those of the author and do not necessarily reflect those of the University of Manchester, CRUK Manchester Institute, Cancer Research UK.

The author declares no competing financial interest.

ACKNOWLEDGMENTS

A.M.J. thanks Rae Lawrence for thoughtful suggestions on the manuscript and technical discussions around the subject.

DOI: 10.1021/acsmedchemlett.8b00500
ACS Med. Chem. Lett. 2018, 9, 1150−1152

A.M.J. is funded by Cancer Research UK (Grant numbers C480/A11411 and C5759/A17098).

REFERENCES

(1) https://en.wikipedia.org/wiki/Luddite (accessed 24th October, 2018).

(2) Klucznik, T.; Mikulak-Klucznik, B.; McCormack, M. P.; Lima, H.; Szymkuć, S.; Bhowmick, M.; Molga, K.; Zhou, Y.; Rickershauser, L.; Gajewska, E. P.; Toutchkine, A.; Dittwald, P.; Startek, M. P.; Kirkovits, G. J.; Roszak, R.; Adamski, A.; Sieredzińska, B.; Mrksich, M.; Trice, S. L. J.; Grzybowski, B. A. Efficient syntheses of diverse, medicinally relevant targets planned by computer and executed in the laboratory. *Chem.* **2018**, *4* (3), 522−532.

(3) Brown, D. G.; Bostrom, J. Analysis of past and present synthetic methodologies on medicinal chemistry: where have all the new reactions gone? *J. Med. Chem.* **2016**, *59* (10), 4443−4458.

(4) Walters, W. P.; Green, J.; Weiss, J. R.; Murcko, M. A. What Do Medicinal Chemists Actually Make? A 50-Year Retrospective. *J. Med. Chem.* **2011**, *54*, 6405−6416.

(5) Roughley, S. D.; Jordan, A. M. The medicinal chemist's toolbox: An analysis of reactions used in the pursuit of drug candidates. *J. Med. Chem.* **2011**, *54*, 3451−3479.

(6) See, for example, the IBM RXN platform; https://rxn.res.ibm.com/ (accessed 24th October 2018).

(7) http://blogs.sciencemag.org/pipeline/ (accessed 24th October 2018.)

(8) *Robotics, Artificial Intelligence, and Medicinal Chemistry: A Look Into the Future*, European Federation of Medicinal Chemistry 24th International Symposium on Medicinal Chemistry, 28th August−1st September 2016, Manchester, U.K.

DOI: 10.1021/acsmedchemlett.8b00500
ACS Med. Chem. Lett. 2018, 9, 1150−1152